First publ. in: EUROVIS 2006: Eurographics/IEEE TCVG Symposium on Visualization; Lisbon, Portugal, May 8th-10th, 2006, pp. 195-202

A Spectral Visualization System for Analyzing Financial Time Series Data

Daniel A. Keim, Tilo Nietzschmann, Norman Schelwies[1], Jörn Schneidewind, Tobias Schreck, Hartmut Ziegler

Department of Computer and Information Science, University of Konstanz, Germany
{keim,nietzsch,schneide,schreck,ziegler}@informatik.uni-konstanz.de

[1] DekaBank, Frankfurt am Main, Germany
Norman.Schelwies@dekabank.de

Abstract

Visual data analysis of time related data sets has attracted much research interest recently, and a number of sophisticated visualization methods have been proposed in the past. In financial analysis, however, the most important and most common visualization techniques for time series data is the traditional line- or bar chart. Although these are intuitive and make it easy to spot the effect of key events on a asset's price, and its return over a given period of time, price charts do not allow the easy perception of relative movements in terms of growth rates, which is the key feature of any price-related time series.
This paper presents a novel Growth Matrix *visualization technique for analyzing assets. It extends the ability of existing chart techniques by not only visualizing asset return rates over fixed time frames, but over the full* spectrum *of all subintervals present in a given time frame, in a single view. At the same time, the technique allows a comparison of subinterval return rates among groups of even a few hundreds of assets. This provides a powerful way for analyzing financial data, since it allows the identification of strong and weak periods of assets as compared to global market characteristics, and thus allows a more encompassing visual classification into "good" and "poor" performers than existing chart techniques. We illustrate the technique by real-world examples showing the abilities of the new approach, and its high relevance for financial analysis tasks.*

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation
H.4 [Information Systems]: H.4 Information Interfaces and Presentation

1. Introduction

Time related data sets are ubiquitous and appear in many application domains in business and science, including finance (stock market data, credit card transactional data), communication (telephone data, signal processing data, network monitoring) or entertainment (music, video). The analysis of time related data sets is a key issue to get insight into the data, to identify patterns, trends or correlations. This knowledge is essential for many data analysis tasks like fraud- and anomaly detection, decision support, prediction or performance analysis. Visualization techniques have been successfully applied to analyze time related data sets in many application scenarios [LKL*04], and a number of sophisticated visualization methods have been proposed in the past [SM03].

In financial analysis, however, the most important and most common visualization techniques for time series data are still charts, including line-, bar-, sequence-, point charts and their variations [Mur99]. Financial analysts use charts almost exclusively to analyze a wide array of assets such as stocks, bonds, futures, commodities, or market indices to forecast future price movements. Price charts for example provide an intuitive graphical representation of an asset's price movement over time by plotting a sequence of prices over a specific time frame, e.g., by using line charts. This graphical representation makes it easy to spot the effect of

key events on a security's price, it's return rates over a period of time and whether it is trading near its highs or near its lows. However, standard price charts do often not allow the easy perception of relative movement in terms of growth rates, which are the key feature of any price-related time series. Additionally, it is a difficult task to compare return rates for groups of assets or asset sectors over time in a single view due to occlusion effects.
To bridge this gap, this paper presents the novel *Growth Matrix* visualization technique for analyzing assets, which extends the ability of existing chart techniques. The *Growth Matrix* not only shows the return rates of assets over fixed time frames, but it also shows their growth rates for each subinterval of the time frame in a single view. At the same time it allows the comparison of the complete set of subinterval growth rates for groups of assets and their relevant indexes, by employing appropriate sorted layouts. Our technique provides a powerful way for analyzing financial data, since it allows the identification of strong and weak periods of assets compared to global market characteristics, and thus allows a more encompassing visual classification into "good" and "poor" assets than existing chart techniques. We provide real world asset market application examples that show the abilities of the new approach and its high relevance for financial analysis tasks.

2. Background

2.1. Financial Analysis Problems

In many financial applications, the two most important asset price series characteristics from an analyst's point of view are *return* and *risk*. Regarding return, analysts and investors are interested in growth rates of an asset price series within certain, often multiple, different time frames. Briefly, growth rate is defined as the ratio between the asset price at the end and the starting point of a time frame interval. Often, such time frames are not fixed per se, but depend on the analysis task at hand. Also, different time intervals may need to be compared simultaneously. Given a price series spanning several years, an analyst may simultaneously consider many interesting subintervals, such as:

- Engagement periods advised by financial consultants
- Short term behavior in reaction to external effects such as natural disasters and policy changes
- Subintervals from investment decisions to the most current price in the time series (current gain/loss of an investment)
- Growth rates spanning long subintervals of fixed length for long term analysis

Regarding the risk criteria, investors may be interested in factors like

- Volatility
- Maximum loss within a given sub time frame
- Longest sub period with continued gains or losses
- Average loss (gain) during a recession (booming) period

Furthermore, important risk features of an asset may be the distribution of the best and worst engagement or disengagement time frames. Note that for the assets series analysis, any of the $\frac{n\times(n-1)}{2}$ existing subintervals in a time series consisting of n elements may be significant. In addition to individually analyzing price series of single assets, comparing the asset with characteristics of similar assets is required. To this end, analysis rely on business sectors which group assets into similar classes (indices). Then, above described measures for an asset are contrasted against the behavior of the whole sector, which can be achieved by normalizing asset prices with mean-aggregated prices of its sector. Based on identification of sector typical time frame characteristics like bullish / bearish market phases, then the question at hand is how the considered asset performs as compared to the sector average. Thereby, individual assets may be benchmarked against certain sector averages. The task is to find visualization techniques that effectively support these analysis tasks.

2.2. Time Series Visualization

To date, the visual analysis of time-related data has received significant research attention in Information Visualization and a number of advanced visualization techniques for various analysis tasks have been proposed. Examples are the Spiral technique [WAM01] to identify periodic patterns in time series data, the Cluster and Calendar based Visualization technique [WS99] to identify patterns and trends on multiple time scales simultaneously, or the TimeSearcher [HS04] to visually explore time related data sets via interactive queries. Due to the increasing volume and complexity of time related data sets, many techniques consider the problem of simultaneously visualizing large as well as long time series data, while maintaining perceptibility of the time- and the value dimension. The Recursive Pattern technique [KAK95] visualizes very large time related data sets by employing pixel based rendering. Other techniques like [BM04, Chu98] use numeric aggregation to scale with the size of time-related data sets. In [HKDS05] the authors generate importance–driven layouts visualizing sets of bar charts, reflecting importance and hierarchical relationships. A nice survey on recent research regarding time series visualization is given in [SM03].

The financial sector is an important domain dealing with complex time dependent data sets. The visual analysis of these kind of data sets is an essential issue in technical finance market analysis to support asset return analysis and decision making processes. In financial analysis, however, the most important and most common visualization techniques for time series data are chart diagrams [Mur99], two-dimensional graphs where the x-axis represents time and the y-axis the dependent variable like the price of an asset. Most common are line- and bar charts, since they provide an intuitive way to get insight into price fluctuations of securi-

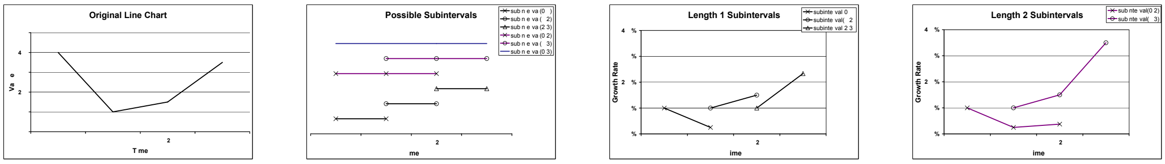

(a) (b) (c) (d)

Figure 1: *Decomposition of a time series into all possible subintervals. (a) Original line chart equivalent with subinterval chart of length* 3*. (b) All possible subinterval time frames. (c,d) Partial line charts of all possible subintervals of length* 1 *and* 2*, normalized with corresponding start values. By displaying normalized subinterval line charts, growth rates are perceivable. Note that for realistic (lager) financial time series, such charts quickly become crowded.*

ties and assets. These charts may be enriched by overlaying aggregate plots, e.g. moving averages. Since there are typically thousands of assets in the market and the analysis of the return rates of each one over time using charts is a difficult task, since numerous charts have to be constructed and compared, there are several approaches that provide compact visual representations of the financial market, e.g. categories/groups of assets, as a whole. Examples are TreeMap approaches [Wat99, Sma05], SOM techniques [DK98] or graph visualization techniques [DG04]. Although these techniques allow compact overviews of the data, they do not allow a detailed analysis or comparison of asset returns over multiple time frames, which is a key issue in financial analysis, and are therefore no real alternative for detailed asset return analysis compared to chart techniques.

By using chart techniques to determine asset growth rates for a certain time frame, the user has to perceive the respective start and end price values, and mentally form the ratio between both. This is not an easy task. Further problems occur when there is high decline in a time series: Reasonable positive or negative growth rates cannot be identified easily. We illustrate this point by an example. Consider the price series $p_0 = 100, p_1 = 60, p_2 = 30, p_3 = 20, p_4 = 22$ it is difficult to determine the growth rate in subinterval $[t_3, t_4]$ as $\frac{22}{20} = 1.1$. This is due to a scaling effect. The absolute difference between p_4 and p_3 (2) is small compared to the absolute difference between t_3 and t_0 (80). As standard line charts are usually linearly scaled for the maximum value, p_3 and p_4 get scaled down quite a bit, making it difficult to read the precise numeric value from the chart. Of course one could easily identify this growth rate if the line chart was normalized to the value of t_3. However, a line chart can be normalized by at most one value, which means this is possible for only the subintervals that have that same starting point, t_3. It is also hard to identify time frames with maximal growth rates, since the analyst has to first determine the growth rates of potentially all subintervals, and then select the best one. However, if not supported by automatic computation, doing this visually on a line chart is difficult as the observer has to memorize many growth rates.

Figure 1 illustrates a simple time series consisting of four values, along with all possible time subintervals. Note that using line charts, it would not be possible to show all subintervals with an individual line chart, as for larger time series the number of subintervals grows too large to display them all in a single display (occlusion and over plotting problems would occur).

3. The Growth Matrix Approach

Based on the preceding analysis, we recognize the need for a technique for simultaneously displaying growth rates for all possible time subintervals in a time series. In Financial Analysis, the *Return Triangle* visualization technique [Deu05, IS05] is a common approach to focus this task. The basic idea is to use a 2D triangular layout with the horizontal and vertical dimensions representing time, in order to visually display asset growth rates on a year-based scale. Color coding is used to identify growth rate quantities. Based on this technique we introduce the pixel-based *Growth Matrix* technique which allows not only the visualization of subintervals at full scale (days, weeks, months), but also allows effective intra- and inter asset analysis. To investigate multiple performance metrics over time the Growth Matrix supports a variety of analysis- and color coding techniques.
The basic idea is to use a 2D layout to place two independent variables, namely, start- end endpoints in time defining respective sub time frames. Mapping the two independent variables onto two different axes in a Cartesian coordinate system, we obtain a reference frame for our two dimensional observation space. We note that by introducing a second independent variable, we loose one (Cartesian) dimension for representing the dependent variable. In terms of the Cartesian coordinate system and the 2D raster pixel display, then only one pixel (rectangle) at the intersection of the respective row and column defined by a subinterval will be available to represent the dependent variable. Note that for one-dimensional bar or line charts displayed in 2D, the second

dimension is free for representation of value magnitude by scaling of bar height or line position. As an alternative, we use an appropriate color-coding scheme to indicate the value at each position in the display. In case of growth rates, we can easily use a color scale ranging e.g., from red (losses) to yellow (side-wards movements) to green (gains).

The *Growth Matrix* can then be defined as a function $GM : (Integer, Integer) \rightarrow Color$ mapping discrete 2D coordinates (i, j) to color codes taken appropriately from a color space. Regarding growth rate normalization, we define two different metrics. The first one is called *growth index*, and is a function $gi^A(i, j)$ based on the relative growth of the price of asset A within the interval $[i, j]$:

$$gi^A(i, j) = \frac{v_j^A}{v_i^A}, \quad i, j \in T, \quad i < j,$$

where i and j are the indices of time stamps corresponding to valid start and endpoints in a global time interval T as defined by the respective time series. v_i^A and v_j^A are the observed prices of asset A at time indices i and j. Then, $gi^A(i, j) \in [0, \infty)$, also referred to as *growth coefficient*, is the relative increase (decrease) of the time dependent variable v in the interval $[i, j]$.
The second growth metric we consider is the *rank index* as a function $ri^A(i, j, DB)$, defined as the $[0, 1]$-normalized rank of a given growth coefficient w r.t. all corresponding growth coefficients in a database of time series DB. Formally:

$$ri^A(i, j, DB) = \frac{1}{|DB| - 1} \sum_{A' \in DB/A} \begin{cases} 1 & gi^{A'}(i, j) < gi^A(i, j) \\ 0 & \text{otherwise} \end{cases}$$

Defining an appropriately normalized color mapping scheme for gi and ri, and plotting the functions in a 2D Cartesian coordinate system with i (j) on the horizontal (vertical) axis, we obtain an upper triangular matrix display of color coded growth rates. The display visualizes the complete spectrum of all possible growth rates in a given time series. Figure 2 illustrates three orientations in the Growth Matrix display which can be interpreted as follows. Along vertical lines, i is kept constant while j is variable. Tracing vertical lines along top direction, we see the returns achieved by investments made at i for increasing holding periods (Figure 2 (a)). Correspondingly, along horizontal lines j is kept constant, while i is variable. Thereby, tracing horizontal lines we can read the return from an investment terminated at a fixed date, considering different initial investment points (Figure 2 (b)). Finally, tracing lines parallel to the matrix diagonal, we can read the returns from investments of constant time periods with variable starting points. In other words, diagonal lines represent sliding windows along the time series, where the window size increases with the lines' distance from the diagonal.

Apart from such local analysis views, Growth Matrices are well suited for evaluating global asset returns, for comparing the return rates of sets of assets, and for identifying

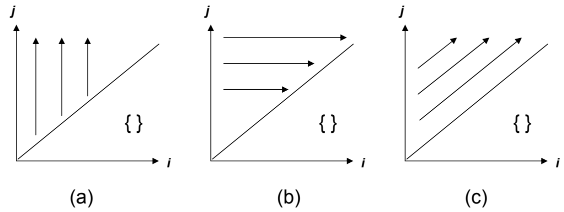

Figure 2: *The* Growth Matrix *coordinate system. i and j denote the start and endpoints of subintervals in the time span given by an original time series. Vertical, horizontal, and lines parallel to the reference frame diagonal can be meaningfully interpreted in terms of hypothetical investment decisions.*

typical growth patterns among sets of time series. In Section 4 we will discuss several applications of the Growth Matrix technique. We note that the Growth Matrix basically is a transformation of time series data with respect to relative growth rates. Up to numerical precision, the original time series can be losslessly reconstructed from the growth rates. Its power for visual analysis lies in the fact that it decomposes an input (one-dimensional) time series into the *spectrum* of all possible subintervals, and visualizes this spectrum. Finally, we note that due to $i < j$, the Growth Matrix occupies just a partition of square display space – the diagonal ($i = j$) and lower triangle ($i > j$) space remains usable for additional views on the given time series, or on appropriately defined database aggregates (see Section 4).
The growth rates of each subinterval in the time series are mapped to the Growth Matrix like illustrated in Figure 3. The resulting visualization allows a detailed analysis of asset return over time. Standard approaches like line charts focus on fixed time frames and force the analyst to construct individual charts for each interval of interest, e.g., for interval $[t_3, t_5]$ in figure 3. Here, asset 1 has a good growth rate in this particular interval $(+8\%)$, but a poor overall growth rate $(-5\%$ in $[t_1, t_5]$). Using the Growth Matrix, the analyst can extract this information from a single view, and at the same time compare individual growth rates over multiple assets. Each matrix can tell a story about the position of the fund in its own peer group. With a peer group of comparable, similar funds, investing in the same asset class, it is easy to find fund managers with persistent achievement. Because of the interval $[0; 1]$ there is no limitation in peer group size. So in a single view you get a comparison of a fund with even a few hundreds others.

4. Application

We have implemented the Growth Matrix technique in an interactive application for visualizing Growth Matrices from time series databases. The system allows interactive selection of input data and application of various preprocessing and benchmarking schemes, supporting the analysis of sin-

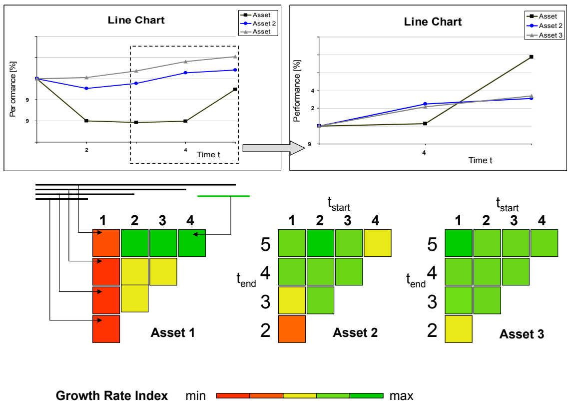

Figure 3: *From Line Charts to Growth Matrix: Simultaneously displaying growth rates for all possible subintervals in one global time series allows to show details on the global and local return rates of assets.*

gle assets as well as contrasting sets of assets. We use a real-world data set containing price series for approximately 10.000 fund assets. In this database, all prices were sampled on a monthly basis during January 1991 and March 2005, leading up to 171 prices per fund (not each asset covered the whole time frame). The database represents European and International funds composed of *bonds* as well as *stock* assets, and also mixed funds. The funds of bonds usually show moderate, rather steady growth patterns, while the stock-based funds usually exhibit higher dynamics (variance) in the price series.

In Section 4.1, we exemplarily analyze each one fund of type bond and stock, identifying typical growth characteristics. In Section 4.2, we extend the basic Growth Matrix by a market-based normalization scheme, allowing to contrast asset-local growth against certain market benchmarks. Section 4.3 then demonstrates the usefulness of the technique for *screening* sets of similar types of funds with respect to their return measured against the market. In Section 4.4, we conclude the application by giving an analytical index for summarization and ranking of Growth Matrices based on a model of investor preferences. Figure 4 details the colormap used in the following discussion.

Figure 4: *Our colormap ranges from shades of red (negative growth) to beige (zero and small growth) through green to black (positive growth). In the following, we apply a linear mapping scheme which caps all growth indices gi at 25% and 250%, respectively. Also, we use the rank index ri for visualizing market-relative rankings. Then, each ri is mapped linearly to the color spectrum.*

4.1. Analysis of Intra-Asset Growth Patterns

Figure 5 (left) shows the Growth Matrix of a fund composed of bonds. As expected, the asset exhibits a rather steady growth pattern across the considered time frame (1991 to 2005). Overall growth was positive throughout all possible investment sub periods, except for some very short term intervals, as the display is dominated by shades of green. The increase in value w r.t. the beginning of potential investment periods is nearly proportional to the investment period length. This is clearly visible, as the intensity of shades of green increases steadily as we consider chart areas more distant from the diagonal. Specifically, there are no periods exhibiting significant negative growth (losses), as opposed to the stock fund example discussed next.

In Figure 5 (right), we illustrate one fund composed of technology stocks. At a glance, one can analyze the impact that the so-called "*new economy*" or "*dot-com*" phenomenon within the technology sector had on this fund. An investor could realize positive, even significant gains by investing into this fund prior to 1997 but selling not later than 2001, which roughly denotes the beginning of the "dot-com" speculative bubble burst. The effects of other investment scenarios can easily be derived from the visualization. E.g., when having invested between 2000 and 2001, one could not have terminated this fund losslessly until the end of the given time frame (March 2005). If we consider long-term investment into this asset held until March 2005 by tracing vertical lines up to the top, we readily perceive that investments made before 1995 did not incur losses even throughout the bubble burst.

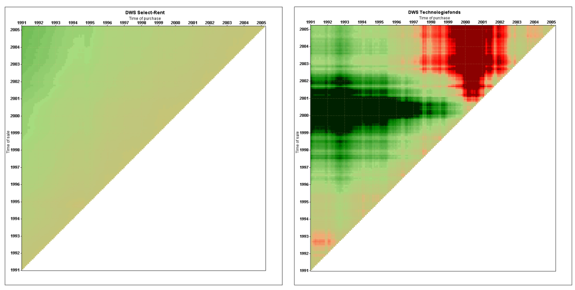

Figure 5: *Growth Matrices (using gi) of a fund composed of bonds (left) and of technology stocks (right). Identifiable are prototypical growth patterns in both kinds of funds. The bond fund achieves moderate, but steady growth throughout all investment subintervals (up to 110 percent gain within the complete time frame of 14 years). Conversely, the stock fund experiences sub periods of significant gains (multiplying in price), but also loosing such intermediate gains in the long run due to the "new-economy" breakdown.*

4.2. Benchmarking Asset Growth Against the Market

The analysis in the preceding Section is an *intra-asset* analysis considering asset-local growth characteristics. We now combine the asset-local analysis with a market-based benchmarking scheme in an integrated Growth Matrix display. This allows not only to analyze isolated asset returns, but also to rate returns in all subintervals against certain market benchmarks. In Figure 6, the Growth Matrices from Figure 5 showing the growth indices *gi* are scaled down and translated to the lower-half (empty) areas within *market-normalized* Growth Matrices over the respective assets. We perform market normalization such that each asset-local growth rate *ri* of a given asset can be compared against the respective growth rates of the other assets in the database. To this end, we find the rank index *ri* of the given asset's growth rates *gi* according to Section 3. The obtained rank index is then linearly mapped to the color scale and visualized. Growth rates of rank higher than the center rank (growth rates higher that the median growth rate) are mapped to increasingly darker shades of green. This in turn implies that such growth rates outperform the market, as represented by its corresponding median growth grates. The opposite applies for growth rates with indices smaller than the center index.

In Figure 6 (left) we compare the bond fund's growth pattern against that of the whole database. We see that the bond fund under performs the market for long investment periods when the investment was done before 1997. This is easily perceivable as the left part of the benchmarked Growth Matrix consists of shades of red indicating below-median returns. This is a typical result as the risk premium of an investment in bonds is usually smaller than it is for an investment in stocks, and our database contains a mix of bond and stock-based fund assets. The opposite is true for the time frame around the "new-economy" bubble burst. During all hypothetical investment periods between 2000 and mid of 2002, this bond-based fund outperformed the market.

In Figure 6 (right) we show the same benchmark for the fund based on technology stocks. Comparing local fund growth rates and the market benchmark, both share structural similarities. We learn that the funds' ranking against the market behaves much like its own local growth pattern: During the new-economy boom, this fund significantly outperformed the market, while after the burst of the speculative bubble, it significantly under performed the market. An exception are investments held between around 1991 and 1994, where the fund under performed the market. Interestingly, we notice a horizontal red "stripe" at the top of the Growth Matrix, indicating that if we held this fund until 2005, we would always have performed worse that the market median, regardless of the point of investment.

We note that here we performed benchmarking of growth rates against our full database of many funds composed of bonds, stocks, or both (mixed asset funds). In practice it is also typical to focus benchmarking on more strict classifications, considering only funds of certain types (stock- or bond-based, regional or sectoral classifications, etc.) for normalization. We point out that already simple, global market normalization can serve to answer intriguing questions such as: *During which periods, if any, does a given fund outperform the market, and by what magnitudes?* Given that an investor can always build a plain portfolio representing the market without incurring the management fees typically associated with funds, this is a central question at the heart of financial analysis.

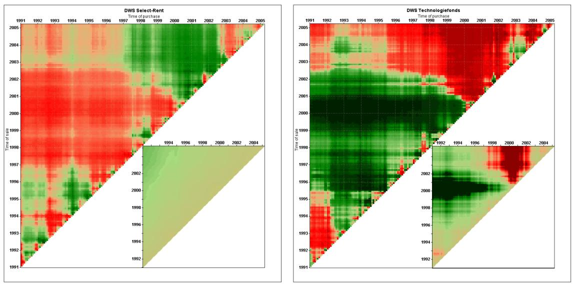

Figure 6: *The bond (left) and stock (right) funds from Figure 5 benchmarked against the market. The large Growth Matrices show the rank index ri for all investment sub periods of these funds, as compared to our full asset database. The small Growth Matrices show the respective asset-local growth patterns (gi) for reference.*

4.3. Comparative Benchmarking of Sets of Assets

Another interesting application of the Growth Matrix lies in *screening* assets from similar classes in order to identify which assets perform better or worse than the average class-specific characteristics in any possible time frame. It is expected that structural patterns of the assets from a given asset class are similar, e.g., funds of European technology stock should exhibit roughly the same high-level growth patterns. Based on simultaneous visualization of assets belonging to a given class, the analyst can then explore for deviations from the general pattern in order to assess discriminative characteristics of individual assets within the given class.

We illustrate such comparative screening by laying out three globally benchmarked Growth Matrices in Figure 7. We chose three funds composed of stocks from Asian-Pacific corporations. The overall pattern is that these funds are dominated by mainly low to negative growth of respective fund prices. Basically, investors would have realized positive returns only when terminating the funds around the year 2000, as the growth rate at this point in time is rather constant and positive, irrespective of the time of initial investment. Among the three stock funds, the rightmost fund performed best, as it shows fewest lossy subintervals in the display.

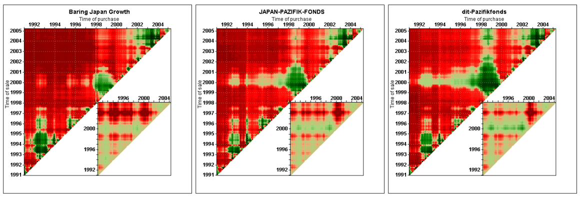

Figure 7: *Three stock-based funds from the Asian-Pacific region. The funds generally under perform the database median, and negative growth patterns are characteristic. Compared to the other funds, the rightmost one performs best.*

Another screening example is given in Figure 8, representing funds composed of technology stocks. While the overall Growth Matrix patterns are similar and reflect the "dot-com" phenomenon, the leftmost fund asset generally exhibits better maximal returns. It has the largest growth rates prior to the technology crisis, and recovers more quickly from it. While the crisis impact is clear, this particular fund has best managed to control losses. Conversely, the rightmost fund performed worst in managing the technology crisis.

4.4. A Visually Motivated Growth Matrix Descriptor

In the preceding discussion, we applied the Growth Matrix technique on several types of assets, and analyzed typical behaviors of growth rates. Intuitively, the visualization communicates performance patterns over time by the distribution of growth rates mapped to color codes. Based on the visualizations, we can go one step further and synthesize an analytical index defined not over raw time series values like is often done in finance, but directly on the Growth Matrix images, trying to capture main visual features. To achieve this, we can model an investor's aggregated preference over all growth rates present in the Growth Matrix. We here specify a class of aggregation functions by means of a weighted sum over all growth indices, considering the *recentness* of each growth index, as well as its respective *period* in a combined weight. We make the following assumptions:

1. Growth index weight increases with recentness. Recentness can be measured by $|T|-j-1$ as the difference between the growth index's last time stamp index j, and the largest time stamp index present in the time series.
2. Growth index weight depends on the time period via a single peaked function centered at a specific time period $\overline{d}$ indicating preferred investment periods.

Formally, we define the *recentness weight function* as $rw(j,|T|) : (Integer, Integer) \rightarrow Real[0,1]$, and the *Time period weight function* as $dw(i,j) : (Integer, Integer) \rightarrow Real[0,1]$. We then calculate the performance index pi for an asset A as the weighted sum over its growth indices $gi^A(i,j)$ (or its rank indices $ri^A(i,j,DB)$), where the combined weight is obtained by multiplying the weights rw and dw:

$$pi^A = \sum_{i=0}^{|T|-2} \sum_{j=i+1}^{|T|-1} rw(j,|T|) \times dw(i,j) \times gi^A(i,j).$$

An instantiation of pi is obtained by specifying rw and dw appropriately. A reasonable specification for a recentness weighting function could be:

$$rw(j,|T|) = \frac{j}{|T|-1} \in [0,1],$$

and also a time period weighting function:

$$dw(i,j) = \max\{1 - s(j-i-\overline{d})^2, 0\} \in [0,1].$$

In these instantiations, rw linearly depends on the recentness of the considered subinterval, and dw quadratically depends on the agreement of the subinterval's period with the preferred time period $\overline{d}$. In the latter function, $s \leq 0$ is a scaling factor for adjusting agreement sensitivity.

The derived performance index pi^A is proposed as a metric for capturing the spectrum of growth characteristics of any asset time series in a single coefficient, according to modeled investor preferences. pi^A is motivated by the Growth Matrix visualization, and may serve as a selection criteria for querying large sets of Growth Matrices for the assets best matching a given investor's preference. Also, it can be used for generating sorted layouts for sets of assets. Additionally, the combined weight $rw \times dw$ could be used for scaling the saturation attribute in Growth Matrix images, visually highlighting subintervals of interest to the investor. Due to space limitation reasons we cannot present applications using this performance index, but will do so in future work.

5. Conclusions

We have introduced the Growth Matrix technique for analyzing time series data by visualizing growth rates defined over all subintervals in the original (raw) time series. We implemented a visualization system based on the Growth Matrix technique, and applied it on a large financial data set consisting of fund price series. We also developed an analytical index to capture key features of the obtained Growth Matrix images. We found the technique very useful for analyzing growth patterns present in asset price time series. In the studied financial context, the Growth Matrix allows the rigorous screening of asset price developments on an individual basis (intra-asset analysis), as well as contrasting many time series simultaneously against each other as well as against the market (inter-asset analysis, and market-benchmark analysis).

Future work will explore the proposed performance index, and extend the system's interactive facilities by appropriately coupling the Growth Matrix display with dynamic line-chart views. This work has focused on growth rates visualization. Regarding the broader field of time series analysis, considering subinterval scores obtained from, e.g., change

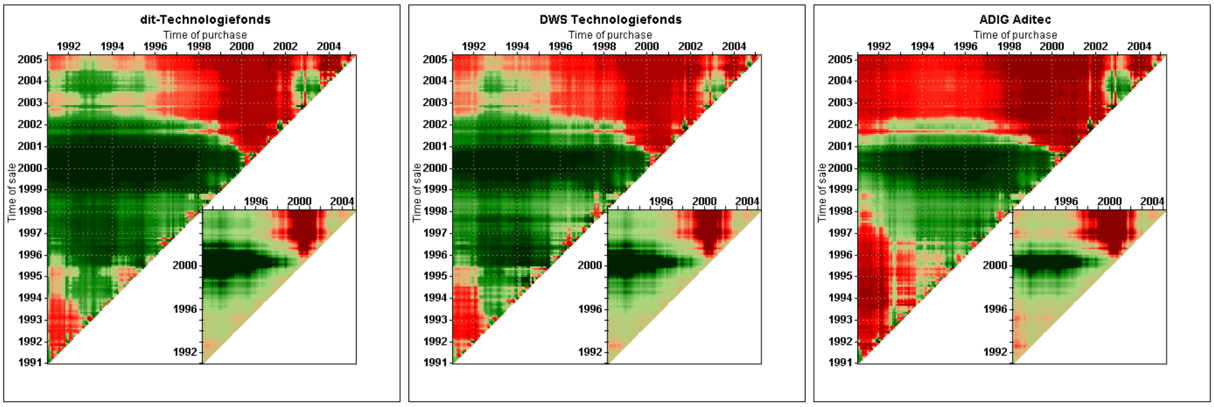

Figure 8: *Three funds composed of technology stocks. By simultaneous visualization of a set of similar assets, these can be readily compared. Strong and weak candidates can be identified.*

point, trend, or anomaly analysis techniques, seems promising for selected applications.

Acknowledgements

Thanks to the DekaBank, Frankfurt am Main, Germany for providing data and valuable consultation. This work was partially funded by the German Research Foundation (DFG) under grant GK-1042 “Explorative Analysis and Visualization of Large Information Spaces”.

References

[BM04] BERRY L., MUNZNER T.: Binx: Dynamic exploration of time series datasets across aggregation levels. In *Poster Compendium to Proceedings of the IEEE Symposium on Information Visualization (INFOVIS’04)* (Washington, DC, USA, 2004), IEEE Computer Society, p. 215.2.

[Chu98] CHUAH M. C.: Dynamic aggregation with circular visual designs. In *INFOVIS ’98 Proceedings of the 1998 IEEE Symposium on Information Visualization* (Washington, DC, USA, 1998), IEEE Computer Society, pp. 35–43.

[Deu05] DEUTSCHES AKTIENINSTITUT E.V.: Return triangles, http://www.dai.de/, 2005.

[DG04] DWYER T., GALLAGHER D. R.: Visualising changes in fund manager holdings in two and a half-dimensions. *Information Visualization 3*, 4 (2004), 227–244.

[DK98] DEBOECK G. J., KOHONEN T. K. (Eds.): *Visual Explorations in Finance with Self Organizing Maps*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1998.

[HKDS05] HAO M., KEIM D., DAYAL U., SCHRECK T.: Importance driven visualization layouts for large time-series data. In *Proceedings of the IEEE Symposium on Information Visualization (InfoVis)* (2005), IEEE Computer Society, pp. 203–210.

[HS04] HOCHHEISER H., SHNEIDERMAN B.: Dynamic query tools for time series data sets: timebox widgets for interactive exploration. *Information Visualization 3*, 1 (2004), 1–18.

[IS05] IBBOTSON STAFF M. W. B. (Ed.): *Stocks, Bonds, Bills, and Inflation 2005 Yearbook Valuation Edition*. Ibbotson Associates, 2005.

[KAK95] KEIM D. A., ANKERST M., KRIEGEL H.-P.: Recursive pattern: A technique for visualizing very large amounts of data. In *VIS ’95 Proceedings of the 6th conference on Visualization ’95* (Washington, DC, USA, 1995), IEEE Computer Society, pp. 279–286.

[LKL*04] LIN J., KEOGH E., LONARDI S., LANKFORD J. P., NYSTROM D. M.: Visually mining and monitoring massive time series. In *KDD ’04 Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (New York, NY, USA, 2004), ACM Press, pp. 460–469.

[Mur99] MURPHY J. J. (Ed.): *Technical Analysis of the Financial Markets A Comprehensive Guide to Trading Methods and Applications*. Prentice Hall Press, 1999.

[SM03] SCHUMANN H., MUELLER W.: Visualization methods for timedependent data – an overview. In *Proc. Winter Simulation Conference* (2003), ACM, pp. 737–745.

[Sma05] SMARTMONEY: Map of the market, http://www.smartmoney.com/marketmap/, 2005.

[WAM01] WEBER M., ALEXA M., MUELLER W.: Visualizing time-series on spirals. In *INFOVIS ’01 Proceedings of the IEEE Symposium on Information Visualization 2001 (INFOVIS’01)* (Washington, DC, USA, 2001), IEEE Computer Society, pp. 7–14.

[Wat99] WATTENBERG M.: Visualizing the stock market. In *CHI ’99 CHI ’99 extended abstracts on Human factors in computing systems* (New York, NY, USA, 1999), ACM Press, pp. 188–189.

[WS99] WIJK J. J. V., SELOW E. R. V.: Cluster and calendar based visualization of time series data. In *INFOVIS ’99 Proceedings of the 1999 IEEE Symposium on Information Visualization* (Washington, DC, USA, 1999), IEEE Computer Society, pp. 4–9.